

REMARKS

Claims 1-11, 13, 14, 30, and 49-70 are pending upon entry of this Amendment. By this Amendment, claims 40-48 are canceled without prejudice; Applicants retain the right to pursue these claims in a Divisional application. Also by this Amendment, new claims 67-70 are added. There are a total of 36 pending claims, four of which are independent.

Amendments

Claims 1-11, 13, 14, 67, and 68 are method claims that generally correspond to Beauregard claims 51-63, 69, 70. Numerous changes are made by the amendments. New claims 67 and 69 reflect features extracted from independent claims 1 and 51, respectively. New claims 68 and 70 clarify that the NIC manager can also handle network traffic for components other than to and from the VMs.

Each of the independent claims is amended to clarify that the data being transferred is an outgoing data frame. Claims 1 and 51 are additionally modified to provide that certain elements of the methods take place within a component of the virtualization software. Various other minor claim restructurings are also implemented. Dependent claims are amended primarily for consistency.

Independent claims 30 and 64 are amended to clarify that the first and second NICs are physical NICs. Claims 30 and 64 are further amended to couple the decision as to whether to transmit the outgoing data frame or discard the outgoing data frame with the determination as to whether the VM that provided the outgoing data frame had higher priority or lower priority. Claims 49, 50, 65 and 66, depending from claims 30 or 64, are likewise amended for consistency.

Previously withdrawn claims 40-48 are canceled without prejudice.

No new matter is entered by this Amendment.

Claim Rejections – Mahalingam and Vega

Claims 1-8, and 10 are rejected under 35 U.S.C. § 103(a) for being unpatentable over U.S. Patent 6,208,616 to Mahalingam et al. (hereinafter referred to as “Mahalingam”) in view of U.S. Patent 7,136,800 issued to Vega (hereinafter referred to as “Vega”).

Claims 1 and 51 are allowable because prior art references Mahalingam and Vega fail to teach or suggest each and every feature set forth in said claims.

For obviousness under 35 U.S.C. § 103(a), each and every claimed feature must be taught or suggested by the prior art reference, or references when combined or modified (MPEP 2143). Applicant need only point out a single feature in each claim that is not disclosed, taught, or suggested by any reference identified in the Office Action to overcome the prior art-based rejection. The following discussion therefore should not be construed as an exhaustive listing of every distinguishing feature set forth in the claims.

The Office Action includes an admission that “Mahalingam does not teach using VM-specific information in the decision making process.” However, the Office Action states that Vega “explicitly teaches using VM-specific information . . . to manage the host computer’s resources” (Office Action, page 4 lines 20-24). Applicants agree with these statements. However, neither Mahalingam nor Vega specifically teach managing NICs using both NIC management information and VM-specific information as set forth in claims 1 and 51.

Specifically, claims 1 and 51 set forth the step of “based on the . . . VM-specific information, selecting a NIC from the plurality of NICs and transferring the outgoing data frame to the computer network over the selected NIC”. While Applicants agree that Mahalingam does discuss selecting a NIC from a plurality of NICs, that decision is based *only* on the NIC management information (e.g., failover status and load balancing information). There is no suggestion in Mahalingam to select a NIC based on VM-specific information or even any information that is analogous thereto. Likewise, Vega does not teach or suggest selecting a resource from a plurality of like resources based on VM-specific information. Rather Vega deals *exclusively* with allocating a percentage (or fraction) of a unitary resource, i.e., processor capacity. Claims 1 and 51 should therefore be allowed because neither Vega nor Mahalingam teach or suggest, “based on the . . . VM-specific information, selecting a NIC from the plurality

of NICs and transferring the outgoing data frame to the computer network” as set forth in claims 1 and 51.

Furthermore, the Office Action provides no rationale to support the modifications to either or both Mahalingam and Vega that would be required to meet the features set forth in the claims discussed above. Thus, the rejections set forth with respect to Mahalingam and Vega fail to satisfy the explicit analysis required to sustain a rejection as set forth in *In re Kahn*, 441 F. 3d 977, 998 (Fed. Cir. 2006) as cited with approval by the United States Supreme Court in *KSR International, Co., v. Teleflex, Inc.*, 127 S.Ct. 1727 (2007).

Claims 1 and 51 are allowable because the prior art lacked motivation to combine and/or modify Mahalingam and Vega in the manner proposed by the Office Action, and to the extent that motivation existed, there would not have been a reasonable expectation for success in the combining of Mahalingam with Vega.

For an obviousness type rejection, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or combine reference teachings (MPEP 2143.01). If motivation is found, then there must still be a reasonable expectation of success in making the combination of references (MPEP 2143.02.I).

Motivation to combine the references is lacking because, for example, Vega is directed to allocating a percentage (or fraction) of a unitary resource, i.e., processor capacity, which is inapplicable to NIC management. The Examiner’s contention that “[a]though Vega is directed to allocating processor time, one skilled in the art understands that a NIC is a computer resource and the same allocation methods can be used” (Office Action, page 5, lines 1-2) is incorrect; the techniques described by Vega for allocating processor resources using VM-specific information cannot be applied to managing a plurality of NICs, as explained in more detail below. Furthermore, Vega is solely directed to managing “processor resources.” There is nothing in Vega, or any other cited reference, that would have led a person of ordinary skill to conceive of applying the techniques of Vega to a NIC management scheme. This is because of the fundamental differences between processor resources and multiple NICs.

For example, processor resources may be allocated among a plurality of VMs simply by making scheduling decisions for the VMs. See, e.g., col. 7, lines 57-61 of Vega. When a VM is

descheduled so that a different VM can execute, it is like hitting a “pause” button, in that there is no activity in that VM during the time it is descheduled. However, when a VM is denied access to a NIC, the VM may still be executing and in fact be generating outgoing data frames, which queue up over time due to the non-availability of the assigned NIC. This causes additional complexities not envisioned by Vega. The present invention describes various mechanisms for dealing with this situation that do not map to any analogous function in Vega. For example, claim 67 (previously part of claim 1) sets forth discarding an outgoing data frame when a NIC is not available. Vega doesn’t contemplate simply discarding instructions from a descheduled VM that cannot be executed, because while outgoing data frames may be discarded without materially affecting the operation of a VM, the instructions must be executed. Since one can’t simply “pause” a stream of outgoing packets like you can a stream of instructions executing on a processor, there would not have been a likelihood of success in applying the teachings of Vega to managing a plurality of NICs unless there is some modification to either or both of Mahalingam and Vega in making the claimed combination. However, there was, at the time the invention was made, no motivation to do so.

In combining Vega with Mahalingam, the Office Action states:

“Motivation to combine these references is common knowledge in the art. Using thread / process specific information (such as priority) is well known when allocating processor time in a multithreaded environment. Logical partitions are a natural extension of a multithreaded operating system, and thus using VM specific information to allocate resources would have been obvious to one of ordinary skill in the art”

(Office Action, page 5, lines 2-7). Applicants respectfully agree that using VM specific information was known to allocate resources as exemplified by Vega. However, this does not satisfy the requirement of a cogent analysis¹ as to how a person having ordinary skill in the art, at the time the invention was made, could have used the teachings of Vega and Mahalingam to arrive at Applicants’ claimed invention, particularly the operation of “based on the . . . VM-specific information, selecting a NIC. . . .”

Applicants respectfully note that in *KSR*, *infra* obviousness was concluded where “all the claimed elements were known in the prior art and one skilled in the art could have combined the

¹ See *In re Kahn*, 441 F.3d 977, 998 (Fed. Cir. 2006) as recently cited with approval by the United States Supreme

elements as claimed by known methods with no change in their respective functions, and the combination would have yielded nothing more than predictable results to one of ordinary skill in the art. *KSR International Co. v. Teleflex Inc.*, 82 USPQ2d 1385, 1395 (2007). In the present case, the elements cannot be combined “with no change in their respective functions.” Once one accepts the notion that the allocation algorithm supplied in Vega cannot be directly used to make a NIC selection due to the disparate natures of the two different types of computer resources, then that person must accept that some adaptation of Vega would be necessary to make management decisions for a plurality of NICs based on VM-specific information. Without such adaptation, there would not have been a reasonable expectation of success. Hence, the facts of the present case cannot be compared with those in *KSR*.

For the reasons set forth above, Applicants respectfully submit that the prior art fails to teach each and every feature set forth in claims 1 and 51, and since there a combination of the references would not have been successful without making significant modifications, there was no suggestion or motivation in the prior art to combine the references. Furthermore, there was no motivation to modify the references as necessary to obtain the features set forth in claims 1 and 51. Since the Office Action provides no such motivation it can be inferred that the Examiner is instead relying on improper hindsight analysis.² Accordingly, Applicants respectfully submit that claims 1 and 51 are allowable over the prior art of record, and early allowance of the same is respectfully requested.

Each of claims 2-11, 13, 14, 52-63, and 67-70 depend from either of claims 1 and 51, and further distinguish from the prior art. With regard to features set forth in claims 67 and 69, (previously residing in claims 1 and 51, respectively) the Office Action states that the argument that neither Mahalingam nor Vega disclose deciding whether to transfer the data “is not persuasive because Mahalingam does disclose this step as a decision state where it is decided whether or not a packet should be transferred . . . and if it is decided not to transfer, the packet is discarded” (Office Action, page 3, lines 2-6; and page 4, lines 18-19). Applicants respectfully point out that claim 67 (and corresponding claim 69) sets forth, “discarding the *outgoing* data

² Court in *KSR International, Co., v. Teleflex, Inc.*, 82 USPQ2d 1385 (2007).

² *In re Kahn*, 441 F. 3d 977, 998 (Fed. Cir. 2006) (“When the Board does not explain the motivation, or the suggestion or teaching, that would have led the skilled artisan at the time of the invention to the claimed combination as a whole, we infer that the Board used hindsight to conclude that the invention was obvious” (citation omitted).)

frame if a decision is made not to transfer the outgoing data frame” (claim 67, line 4; emphasis added) whereas Mahalingam discloses discarding only *incoming* data frames. Since Mahalingam does not teach discarding outgoing data frames, Applicants respectfully submit that claims 67 and 69 are patentable over Mahalingam and cited references.

Claim Rejections – Macchiano

Claim 30 is rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent 7,111,303 issued to Macchiano et al. (hereinafter referred to as “Macchiano”) in view of Vega and Mahalingam. Claim 64 is a “Beauregard” style claim that generally corresponds in scope to the claim 30, which is a method claim. Claims 49 and 50 depend from claim 30, whereas corresponding claims 65 and 66 depend from claim 64. Applicants note that the Office Action appears to have swept the claims 49, 50, 65, and 66 in with sets of claims corresponding to earlier claims (see, e.g., Office Action, page 11, last three lines).

The prior art fails to teach or suggest each and every feature set forth in claims 30 and 64. For example, the Office Action states, “Macchiano also discloses, ‘the virtual computer system also comprising a first physical network interface card (NIC) and a second physical NIC for connecting to the computer network’ as describing each user portion having a virtual NIC and the computer system may also contain multiple physical NICs (see col. 3, lines 56-58 and Fig. 1 comp. 42, 44; col. 5, lines 4-6)” (Office Action page 10, lines 12-16). Applicants have carefully reviewed the indicated portions of Macchiano and find no support there for a plurality of *physical* NICs. Macchiano does describe (referring to Macchiano, Fig. 1) a computer system 10 having two guest VMs 12, 14, each having a corresponding *virtual* NIC 42, 44. The indicated portions of the specification of Macchiano likewise only show virtual NICs, not physical NICs. To further clarify this distinction, each recitation of NIC in claims 30, 49, 50, 64, 65, and 66 is amended to recite “physical NIC.”

Furthermore, Macchiano does not disclose “transferring the outgoing data frame over an available one of the physical NICs if the one of the first VMs having higher priority provided the outgoing data frame; or discarding the outgoing data frame if the other of the VMs provided the outgoing data frame” as now set forth in claims 30 and 64. Furthermore, none of the remaining cited references teach this feature. Applicants respectfully submit that claims 30 and 64, and

claims depending therefrom, are allowable and therefore request reconsideration in view of the amendments now made to these claims.

Additional dependent claims not specifically addressed above should be allowed for at least the reasons mentioned above with respect to claims from which the additional dependent claims depend. Accordingly, Applicants respectfully submit that the present Application is in condition for allowance. Applicants respectfully request reconsideration of the outstanding rejections in light of the above and a Notice of Allowance. The Examiner is invited to contact the undersigned at 650-427-2390 to discuss any additional changes the Examiner may feel is necessary in light of this Amendment.

Date: July 3, 2008

Respectfully submitted,

/Leonard Heyman/
Leonard Heyman
Reg. No. 40,418
Attorney for the Applicant

APPENDIX – CLAIMS AS AMENDED

For the Examiner's convenience, the claims are presented below as amended:

1. A method for responding to a request to transfer an outgoing data frame from a virtual computer system to a computer network, the outgoing data frame comprising at least data to be transmitted and at least one of a layer 2 and layer 3 destination address, the virtual computer system comprising one or more virtual machines (VMs) executing on a host platform via virtualization software, the virtualization software comprising one or more layers of software interfacing between the VMs and the host platform, the virtual computer system further comprising a plurality of physical network interface cards (NICs), the method comprising the steps of:

obtaining access by a NIC manager to the outgoing data frame, the outgoing data frame being provided by one of the VMs, the NIC manager being a component of the virtualization software;

receiving, in the NIC manager, NIC management information related to one or more of the plurality of NICs;

receiving, in the NIC manager, VM-specific information related to one or more of the VMs in the virtual computer system; and

based on the NIC management information and the VM-specific information, selecting a NIC from the plurality of NICs and transferring the outgoing data frame to the computer network over the selected NIC.

2. The method of claim 1, in which the VM-specific information indicates an amount of network bandwidth that is allocated to the one of the VMs that provided the outgoing data frame.

3. The method of claim 67, in which the VM-specific information indicates an allocation to the one of the VMs that provided the outgoing data frame of an amount of network bandwidth and a the decision is made not to transfer the outgoing data frame when transferring the outgoing data frame would cause the allocation of network bandwidth to be exceeded.

4. The method of claim 1, in which the VM-specific information indicates a priority of the one VM relative to the priorities of other ones of the VMs.

5. The method of claim 1, in which the NIC management information indicates which one or more of the plurality of NICs is available for the transfer of the outgoing data frame.

6. The method of claim 5, in which the NIC management information further indicates a pending data transfer load for each of the plurality of NICs that are available for the transfer of the outgoing data frame.

7. The method of claim 1, in which a load distribution function is used in performing the selecting of the NIC over which to transfer the outgoing data frame.

8. The method of claim 7, wherein:

the one or more VMs comprises at least a first VM and a second VM and the plurality of NICs comprises at least a first NIC and a second NIC; and

the load distribution function substantially always routes outgoing data frames provided by a first VM over a first NIC as long as the first NIC is available, and substantially always routes outgoing data frames provided by a second VM over a second NIC as long as the second NIC is available, and routes outgoing data frames provided by the first VM over the second NIC

if the first NIC is not available, and routes outgoing data frames provided by the second VM over the first NIC if the second NIC is not available.

9. The method of claim 8, in which the outgoing data frames provided by the first VM are distinguished from outgoing data frames provided by the second VM by reference to a source physical address contained in a header of each outgoing data frame.

10. The method of claim 1, in which the NIC management information indicates whether a failover is occurring on one of the NICs.

11. The method of claim 10, in which the one VM is temporarily suspended if a failover is occurring on the one of the NICs.

12. (Canceled.)

13. The method of claim 67, wherein, if a decision is made not to transfer the data, a further decision is made whether to suspend the one VM.

14. The method of claim 1, wherein, if a decision is made not to transfer the data, a further decision is made whether to migrate the one VM to another computer system.

Claims 15-29 (Canceled.)

30. A method for responding to requests to transfer outgoing data frames from a virtual computer system to a physical computer network, the virtual computer system comprising a first VM and a second VM, the virtual computer system also comprising a first physical

network interface card (NIC) and a second physical NIC for connecting to the computer network, the method comprising the steps of:

determining that one of the first and second physical NICs is not available for transferring data;

determining that one of the first and second VMs has a higher priority than an other of the first and second VMs; and

for each of the outgoing data frames:

determining whether the first VM or the second VM provided the outgoing data frame; and

transferring the outgoing data frame over an available one of physical NICs if the one of the VMs having higher priority provided the outgoing data frame; and

discarding the outgoing data frame if the other of the VMs provided the outgoing data frame.

Claims 31-48 (Canceled)

49. The method of claim 30, further comprising suspending the other of the VMs in response to determining that the one physical NIC is not available.

50. The method of claim 30, further comprising migrating the other of the VMs to another computer system in response to determining that the one physical NIC is not available.

51. A computer program embodied in a tangible, computer-readable medium, the computer program performing a method for responding to a request to transfer an outgoing data frame from a virtual computer system to a computer network, the outgoing data frame comprising at least data to be transmitted and at least one of a layer 2 and layer 3 destination address, the virtual computer system comprising one or more virtual machines (VMs) executing

on a host platform via virtualization software, the virtualization software comprising one or more layers of software interfacing between the VMs and the host platform, the virtual computer system further comprising a plurality of physical network interface cards (NICs), the method comprising the steps of:

obtaining access by a NIC manager to the outgoing data frame, the outgoing data frame being provided by one of the VMs, the NIC manager being a component of the virtualization software;

receiving, in the NIC manager, NIC management information related to one or more of the plurality of NICs;

receiving, in the NIC manager, VM-specific information related to one or more of the VMs in the virtual computer system;

based on the NIC management information and the VM-specific information, selecting a NIC from the plurality of NICs and transferring the outgoing data frame to the computer network over the selected NIC.

52. The computer program of claim 51, in which the VM-specific information indicates an amount of network bandwidth that is allocated to one of the VMs that provided the outgoing data frame.

53. The computer program of claim 69, in which the VM-specific information indicates an allocation to the one of the VMs that provided the outgoing data frame of an amount of network bandwidth and the decision is made not to transfer the outgoing data frame when transferring the outgoing data frame would cause the allocation of network bandwidth to be exceeded.

54. The computer program of claim 51, in which the VM-specific information indicates a priority of the one VM relative to priorities of other ones of the VMs.

55. The computer program of claim 51, in which the NIC management information indicates which one or more of the available plurality of NICs is available for the transfer of the outgoing data frame.

56. The computer program of claim 55, in which the NIC management information further indicates a pending data transfer load for each of the plurality of NICs that are available for the transfer of the outgoing data frame.

57. The computer program of claim 51, in which a load distribution function is used in performing the selecting of the NIC over which to transfer the outgoing data frame.

58. The computer program of claim 57, wherein:
the one or more VMs comprises at least a first VM and a second VM and the plurality of NICs comprises at least a first NIC and a second NIC; and
the load distribution function:

substantially always routes outgoing data frames provided by a first VM over a first NIC as long as the first NIC is available;

substantially always routes outgoing data frames provided by a second VM over a second NIC as long as the second NIC is available;

routes outgoing data frames provided by the first VM over the second NIC if the first NIC is not available; and

routes outgoing data frames provided by the second VM over the first NIC if the second NIC is not available.

59. The computer program of claim 58, in which the outgoing data frames provided by the first VM are distinguished from outgoing data frames provided by the second VM by referring to a source physical address contained in a header of each outgoing data frame.

60. The computer program of claim 51, in which the NIC management information indicates whether a failover is occurring on one of the NICs.

61. The computer program of claim 60, in which the one VM is temporarily suspended if a failover is occurring on the one of the NICs.

62. The computer program of claim 69, wherein, if a decision is made to not transfer the outgoing data frame, a further decision is made whether to suspend the one VM.

63. The computer program of claim 69, wherein, if a decision is made to not transfer the outgoing data frame, a further decision is made whether to migrate the one VM to another computer system.

64. A computer program embodied in a tangible, computer-readable medium, the computer program performing a method for responding to requests to transfer outgoing data frames from a virtual computer system to a physical computer network, the virtual computer system comprising a first VM and a second VM, the virtual computer system also comprising a first physical network interface card (NIC) and a second physical NIC for connecting to the computer network, the method comprising:

determining that one of the first and second physical NICs is not available for transferring data;

determining that one of the first and second VMs has a higher priority than an other of the first and second VMs; and

for each of the outgoing data frames:

determining whether the first VM or the second VM provided the outgoing data frame; and

transferring the outgoing data frame over an available one of the physical NICs if the one of the VMs having higher priority provided the outgoing data frame; or

discarding the outgoing data frame if the other of the VMs provided the outgoing data frame.

65. The computer program of claim 64, wherein the method further comprises suspending the other of the VMs in response to determining that the one physical NIC is not available.

66. The computer program of claim 64, wherein the method further comprises migrating the other of the VMs to another computer system in response to determining that the one physical NIC is not available.

67. The method of claim 1, further comprising:
deciding, based on the NIC management information and the VM-specific information, whether to transfer the outgoing data frame;
discarding the outgoing data frame if a decision is made not to transfer the outgoing data frame; and
performing the transferring of the outgoing data frame only if a decision is made to transfer the outgoing data frame.

68. The method of claim 1, wherein the NIC manager is additionally provided with access to outgoing network frames from components of the virtual computer system other than the VMs.

69. The computer program of claim 51, wherein the method further comprises:

deciding, based on the NIC management information and the VM-specific information, whether to transfer the outgoing data frame;

discarding the outgoing data frame if a decision is made not to transfer the outgoing data frame; and

performing the transferring of the outgoing data frame only if a decision is made to transfer the outgoing data frame.

70. The method of claim 51, wherein the NIC manager is additionally provided with access to outgoing network frames from components of the virtual computer system other than the VMs.